

Quantum Individuality

A Senior Honors Thesis

Presented in Partial Fulfillment of the Requirements for graduation
with research distinction in Philosophy in the undergraduate colleges of
The Ohio State University

by
Dan Giglio

The Ohio State University
May 2011

Project Advisor: Professor Robert Kraut, Department of Philosophy

Contents

1. Introduction	3
2. Individuality	4
3. Indiscernibility	7
4. Four Arguments for Non-individuality	16
5. Metaphysical Implications of Non-individuality	30
6. Conclusion	35
7. References	36

1

Introduction

This essay addresses a number of thorny issues which all arise from a single claim: quantum particles are (in some sense of the word) not *individuals*. This claim is commonplace to the physicist, puzzling to the philosopher, and simply nonsensical to the layman. Accordingly, the task that underpins this entire essay is to lead the way from the “nonsensicality” knee-jerk reaction to a precise understanding of why it must be the case that quanta cannot be regarded as individuals. Additionally, I will argue that these non-individuals, as such, constitute a counterexample to Leibniz’s Law (LL). The thesis of this essay is therefore twofold.

One motif that will emerge throughout the course of this essay is that the concepts of identity, individuality, and indistinguishability – insofar as they are employed in quantum contexts – are inextricably tied together into one large, knotted mess. I am not certain that this knot can ever be fully untied. I am especially compelled to think that previous attempts to understand these concepts *in isolation from each other* have been a bit like tugging forcefully at both ends of the shoelace without realizing that they are parts of the same whole; this only serves to tighten the knot. I can only hope that my discussion will loosen the knot somewhat, so that we can at the very least – and still with a great deal of effort – slip off our shoes.

2

Individuality

It goes without saying that the advent of quantum mechanics (as well as special and general theories of relativity) profoundly shifted the worldviews of early 20th century scientists and laymen alike. Although it was nearly a century ago that these radical theories stomped their way through the pond of science – mixing up mud and sediment – things have not quite returned to the way they once were, and the water is still quite murky. As Feynman once famously remarked, “I think I can safely say that nobody understands quantum mechanics.”¹

According to Putnam, the reason that scientists were so shocked 100 years ago was not that science was wrong in the *details*, but in the *big picture*.² The Einsteinian revolution pulled the rug out from beneath our feet, so to speak. If science had been slightly wrong about the value of a fundamental constant – e.g. the charge of the electron, or the distance between the earth and the sun – scientists would have been able to cope relatively easily with the changes in their textbooks. However, the areas in which Newtonian physics proved to be wrong were of a much grander scale. The basic metaphysical picture of the world that existed before Einstein, Schrodinger, and Heisenberg had been completely dismantled, a profoundly unintuitive one having taken its place. No longer can matter be regarded merely as “bodies in motion”; certain phenomena can only be properly explained if we endow matter with certain wavelike properties. Notions of absolute position, absolute momentum, and absolute simultaneity were deemed

¹ *The Character of Physical Law* (BBC, 1965), ch. 6.

² This was stated by Putnam in a televised interview with Bryan Magee which can be found here: <http://www.youtube.com/watch?v=cG3sfrK5B4E>. The remainder of this paragraph is more or less a summary of one part of that interview.

illusory. In this section I will focus on one particularly upsetting consequence of the new quantum theory: the evaporation of the concept of individuality.

Now what exactly is meant when we say of an object that it is an individual? In legal contexts, this might mean that the object (a person, or perhaps a corporation) has certain rights which need to be recognized and respected. In psychological contexts, this might mean that the object (a person) exhibits some sort of self-awareness or autonomy. Sometimes we simply use the term ‘individual’ interchangeably with an age-old conception of what physical objects are: physical objects (or, individuals) are those things whose parts move in concert with each other when the object (or, individual) is moved. However, I will not think of individuality in terms of objecthood in this essay. Instead, I will construe individuality in terms of *distinguishability*: something is an individual insofar as it can be *distinguished* or *discerned* from its surroundings, or from its neighbors.

If we accept this construal of individuality, then what can it possibly mean to say that an object is a *nonindividual*? Paul Teller provides a beautiful example (we will return to this example throughout the course of this essay):

Suppose that on Monday I put a silver dollar into my piggy bank, and on Tuesday I put in a second ... On Wednesday I pull one of the silver dollars out and ask, is this the first or the second that I deposited yesterday? Even supposing these two coins to be qualitatively identical, we take there to be a fact of the matter ... even though I may not be in a position to determine which coin is which.

Contrast the last case with one in which on Monday I make a deposit of one dollar in my checking account, and on Tuesday I make a second one dollar deposit. On Wednesday, I appear at the bank wanting to withdraw one dollar. But the teller will not be able to make anything of it if I insist: “And be sure the dollar you give me is the one I deposited on Monday, not the one I deposited on Tuesday!”³

³ Teller (1998), p. 114.

The reason that the teller (lowercase ‘t’) cannot make any sense out of Teller’s (uppercase ‘T’) request is that the dollars in his bank account are not individuals; not a single one of them can be distinguished from the bunch, in contrast to the silver dollars in the piggy bank. Analogously, I will be claiming in this essay that quanta are non-individuals: it only makes sense to speak of *collections* of electrons, and never a *particular* electron. As David Griffiths puts it, “there is *no such thing* as “this” electron or “that” electron; all we can legitimately speak about is “an” electron.”⁴

But much work remains to be done before we can demonstrate that the nature of quantum particles is indicative of non-individuality, and furthermore that such particles constitute a counterexample to LL. Before reaching this conclusion, we first must achieve a very firm grasp on the notion of indiscernibility present in LL. From this point, I will explain why in a number of circumstances quantum particles ought to be regarded as non-individuals, just like the dollars in Teller’s bank account. In each case, I will argue that LL fails *precisely because* of the non-individuality of quanta.

⁴ Griffiths (2005), p. 204.

3

Indiscernibility

3.1 *Leibniz's Law*

Although the notion of indiscernibility may strike the reader as fairly straightforward, I will argue in this section that it desperately stands in need of revision. Short of these revisions, I do not believe that LL can be discussed as precisely as it needs to be in this essay.

Now there are many ways to express LL. The formulation closest to Leibniz's original wording would resemble the following: two things are the same when they can be substituted for one another while preserving truth. In a more modernized form, the principle is expressed through talk of properties: if two things have all their properties in common, then they are one and the same thing. To state this in the language of second order logic:

$\rightarrow =)$ where 'x' and 'y' refer to specific objects and 'F' refers to a property instantiated by those objects. But it is by no means clear at this point how to unpack the notion of 'property' employed in the above formulation: which properties do we deem to be capable of grounding an act of discernment? In other words: what properties are allowed into the extension of F?

French and Krause suggest three ways to interpret the extension of the term 'properties' in the above principle.⁵ The weakest reading liberally attributes the widest scope to this term: each and every property is eligible to serve as a distinguishing property between two objects. If we assume for the moment, quite controversially, that self-identity is a genuine property, then the act of discernment between the two objects x and y becomes trivial. This is because x clearly

⁵ French and Krause (2006), p. 10.

has at least two properties which *y* does not: being identical to *x*, and being different from *y*. These are two properties that *y* cannot have; for if it did have these properties, then we would be speaking of only one object instead of the two that we stipulated.

Allowing self-identity into the extension of *F* is clearly problematic. This is because such an interpretation does not accurately reflect the way in which we as humans go about performing acts of discernment. We do not examine objects to ascertain whether or not they are self identical. We investigate physical properties like position, momentum, and spin. Outside the laboratory, we investigate secondary qualities like how the objects look, sound, feel, etc. If the information is available, we can even examine the historical properties of the objects (where it has been, which of its neighbors it has interacted with, etc.) The list goes on, but it does not include the property of self-identity. For this reason, French and Krause, by default, leave this problematic property outside the extension of *F*.

But we can be even more stringent. The second interpretation of *F*'s scope excludes spatio-temporal properties. Why exclude such properties? The explanation for such a maneuver lies in the fact that spatio-temporal properties often take the form of *relations* among objects. A baseball cannot travel at 10 m/s *simpliciter*; it must travel at 10 m/s relative to some other object or set of coordinate axes. Some philosophers believe that in order to stand in such a relation, we presuppose the numerical diversity of the very object in question. However, this point is quite contentious. It not clear at all what the connection is between that participating in a relation (such as traveling at 10 m/s relative to home plate) and being numerically diverse; perhaps there are a dozen other balls traveling at exactly the same speed!

But we can tighten the screws even more than this. The third and last interpretation fills *F*'s extension with only monadic, non-relational properties. A statement of LL in these terms is a

strong statement indeed: the only properties that can serve to individuate x from other objects are those that can be stated without reference to anything but x . Under this regime, a property such as *being two feet tall* is kosher, whereas a property such as *being two feet from the leaning tower of Pisa* is not.

I am not partial to any of these suggestions. The first (include *all* properties) is too liberal in *any* context, as we have already seen. Furthermore, the second suggestion excludes the very properties that we are interested in: the spatio-temporal properties (this is, after all, an essay which addresses issues regarding the position, momentum, etc. of quantum particles). Lastly, the third suggestion (include only monadic, non-relational properties) is too narrow in quantum mechanical contexts. This is because all quantum particles of a given type share exactly the same monadic properties (mass, charge, etc.) with one another. Therefore, for the remainder of this essay we will adopt a fairly wide scope for F . Spatio-temporal properties, as well as monadic properties, will be assumed to be capable (at least in principle) of grounding an act of discernment. But self-identity will just have to watch the game from the sidelines.

3.2 *What is indiscernibility?*

Once we have decided upon what properties are to be allowed in an act of discernment, a semantic question must be answered before we are to have a firm grasp of LL: what is meant by “indiscernibility?” The notion of indiscernibility has occupied a key role in our efforts to clarify another, and perhaps more puzzling, notion: identity. But that is not to say that the notion of indiscernibility is merely instrumental to philosophizing about identity; it is rich enough to be examined as an end in itself, and not just a means. I wish to suggest a few refinements that bare-bones indiscernibility is desperately in need of.

To see the import of the concept of indiscernibility, we first turn to Quine. According to Quine, in order to fully master count nouns such as ‘apple,’ ‘hammer,’ or ‘bunny,’ it is insufficient to “be able to decide whether a portion of space-time is taken up with apples; we must also be able to determine where one apple leaves off and an adjacent one sets in.”⁶ In other words, language users must be able to perform an act of discernment between groupings of tokens of the same type. Only after they have accomplished this feat for a number of lexical items – and a number of times for each lexical item – are we prepared to call them *competent* language users. Indiscernibility is therefore an obstacle that Quine sets in place for every speaker to overcome as he/she strives to grasp the meaning of certain terms. Flowing from this is the idea that indiscernibility is also an obstacle to the act of counting: before I can begin to tell you how many potatoes are in the pantry, I must first acquire the ability to discern one spud from another.

So far, indiscernibility does not seem to present us with any deep, perplexing problems. Two things are indiscernible when they cannot be told apart from each other. Plain and simple. But there are certainly cases in which it is reasonable to ask “indiscernible to whom?” For example, I am simply not able to detect a difference between the 1970 version of *Let it Be* and the remastered 2003 version, *Let it Be...Naked*. But to a recording expert, these are horses of a different color. So perhaps indiscernibility *simpliciter* is not a fully coherent concept.

In effect, this forces us to construe indiscernibility as a three-place relation rather than a two-place relation. Its inputs include not only the two things that are indiscernible from each other, but also an agent that judges them to be so.⁷ It is also worth noting that by including an

⁶ Quine (1949).

⁷ This is completely analogous to the move that Russell makes in formulating his Multiple Relation Theory of Judgment. Noticing that it is problematic to construe the belief relation as one that obtains between a believer and an object of belief, he suggests that belief is a relation among several objects. For example, Othello’s belief that

agent as a component of the indiscernibility relation, we have transformed indiscernibility from an absolute notion to a relative one. Indiscernibility is not something that is objectively “out there” in the world; there is a subjective character inherent to any claim of indiscernibility.

There is a second aspect in which our everyday construal of indiscernibility must be relativized. Imagine that – as is often the case – a student is forced to purchase the latest edition of a textbook which features incredibly minimal changes from the old (and less expensive) edition. For the simplicity of the example, let us also imagine that the only difference between the old and new versions consists in the correction of typographical errors. Now upon comparing the old and new versions, it would be apt for the student to say something along the lines of “these two books are *indiscernible* from each other!” Needless to say, the student realizes that there are numerous properties that the two books do not have in common. One is being held in his left hand, the other in his right; one was published this year, the other was published last year. This does not, however, reduce the student’s complaint to a mere colloquial expression or piece of rhetoric. There is a very real sense in which the two books are indiscernible from each other: the *content* of the old book is indiscernible from the *content* of the new book.

It seems that what is going on here is that indiscernibility has been relativized to a category of properties. In the above example, the student’s indiscernibility claim held with respect to what can perhaps be called the “pedagogical properties” of the textbooks. Because the books share identical *content* a student would be able to undergo exactly the same learning-experience from both books; they are pedagogically indiscernible from each other.

3.3 Epistemology or Metaphysics?

Desdemona Loves Cassio is *not* a relation between Othello and a single object, Desdemona’s love for Cassio. Instead, “the relation called ‘believing’ is knitting together into one complex whole the four terms Othello, Desdemona, loving, and Cassio” (1912, p. 126).

It is not immediately clear at this point whether we should consider indiscernibility to be an epistemological notion or a metaphysical one. In the case of the former, two objects are indiscernible in virtue of the fact that the agent/judge is *ignorant* to any individuating properties, regardless of whether or not there exist any. In the case of the latter, two objects are indiscernible because there do not exist any properties in the world which can serve to individuate them, regardless of the agent's knowledge (or his very existence, for that matter). Paul Teller creates an example that illustrates this concern well.⁸ Imagine that two friends hold the ends of a rope. Each jerks his hand at the same time, thus creating two waves in the rope – one at either end – both traveling to the center of the rope. When these two waves meet at the center, they merge into a single bump. They then pass through each other and continue down the length of the rope.

Why is it that the two waves are indiscernible at the moment when they pass each other? Is it because they cannot be told apart by any human scientist; or is it simply because there do not exist any properties which serve to differentiate the two waves? In the case of the former, the indiscernibility in question is of the epistemological variety; in the case of the later, the metaphysical variety.

At first blush, it would seem that we are dealing with merely epistemic indiscernibility. One of the waves has the historical property of having originated from Person A's hand (and likewise for the wave Person B created), and although we cannot "see" this property like we can see the position of each wave, it is there nonetheless, and it serves to make the two wave distinct from each other. Commonsense metaphysical realism regarding properties leads us to this conclusion.

⁸ Teller (1983).

However, one may be partial to the idea that metaphysical questions “may very well depend, among other things, on what *exemplified* [italics are mine] characteristics there are which could at least in principle serve to individuate.”⁹ And because historical properties are not *exemplified* in the same way that physical properties are, it follows that they do not have the power to distinguish the two waves from each other. So it is not merely the case that we are ignorant to the key individuating property because we lack the right technological devices to measure historical properties; not even God could differentiate the two waves because at this point in time there simply do not exist any *exemplified* properties capable of differentiation. It is important to note that his metaphysical claim is much stronger than the epistemological claim.

The two-waves example is helpful in framing the question at hand, but at this point we seem to have reached an impasse. Short of deciding whether or not historical properties are genuine properties when it comes to an act of discernment, we will not be able to reach a conclusion regarding the epistemic/metaphysical nature of indiscernibility.

3.4 Problems in the Laboratory

Imagine a two-particle system of electrons. In order to measure the physical properties of either electron, we must to interact with it in some way. For example, we can bombard it with a stream of photons. However, such an interaction with this system will “inevitably and unpredictably alter its state, raising doubts as to whether the two had perhaps switched places.”¹⁰ This predicament has been dubbed the “problem of measurement,” and the experiment just mentioned is only one particular manifestation of the problem. For example, imagine we wish to measure the temperature in the room. If we wait a suitable amount of time after placing our thermometer in the center of the room, the alcohol in the bulb will reach thermal equilibrium

⁹ Teller (1983), p. 310.

¹⁰ Griffiths (2005), p. 204.

with the room. If the room is hotter than the thermometer was initially, then heat will flow into the alcohol causing it to expand; if the room is colder than the thermometer was initially, then heat will flow out of the alcohol causing it to contract. In either case, the length of the thermometer is calibrated in such a way that relates the volume of the liquid to a reading of the room's temperature. But in order for such expansions/contractions to occur, the thermometer must either steal or add heat to the room, thus altering the room's temperature. The very act of measurement *changes* the property we are interested in measuring! A similar situation arises when one tries to measure the air pressure in a tire. In order for the tiny piston on the gauge to shoot outward it must steal some air from the tire, thus changing the pressure inside the tire.

It is probably not clear at this point why state-modification is in any way detrimental to the task of discernment. As long as we glean physical information from one of the two electrons, it might be said, it would make no difference what final state the system finds itself in. This line of reasoning is not quite right, however. To show why state-modification *really is* detrimental to the task of discernment, it will help to imagine a macroscopic manifestation of the problem of measurement. For example, imagine that my task is to discern two beverages from each other. Perhaps I am participating in a blindfolded Coke-Pepsi taste test. Just as the two electrons each have a unique position (if only briefly, and only immediately after an act of measurement) that would serve to distinguish them, the two beverages each have a unique flavor that serves to distinguish them. But let us imagine that quantum phenomena have emerged at the macro-level, and that by sipping the beverage in front of me, I *change its very flavor!* This is analogous to the fact that measuring an electron's position serves to change its very position.¹¹ If such phenomena occurred at the macro-level, then it would be impossible to stage such a taste test; the answer key that states which beverage was Coke and which was Pepsi becomes obsolete after I

¹¹ More accurately: the act of measurement forced it to "take a stand" (David Griffiths' lovely phrase).

take a sip. Just as a measurement on the two-electron system “inevitably and unpredictably alter[s] its state, raising doubts as to whether the two had perhaps switched places,”¹² we could not be know if my act of sipping caused the flavor to change. So we see that in order to discern object A from object B, the system that contains A and B must remain unaltered by the act of discernment. If the final state happens to be altered as a result of the act if discernment, then we haven’t really discerned anything, since we can’t be sure which object is *now* A and which object is *now* B.

¹² Griffiths (2005), p. 204.

4

Four Arguments for Non-Individuality

4.1 *First Argument: Quantum Statistics*

We are now equipped well enough to motivate the first half of the thesis – that quantum particles are not individuals in any meaningful sense of the word. Imagine that our system consists of two particles, call them ‘A’ and ‘B’, which can be distributed across two different energy states, call them ‘ ϵ_1 ’ and ‘ ϵ_2 ’. There would seem to be four possible states for the overall system:

- 1) Both A and B in ϵ_1
- 2) Both A and B in ϵ_2
- 3) A in ϵ_1 and B in ϵ_2
- 4) A in ϵ_2 and B in ϵ_1

Let us assume that each configuration is equiprobable. Notice that the only difference between (3) and (4) is the exchange of A and B in their respective energy levels. Such a swap would be undetectable by us in principle, which means that for the purposes of experimentation, we can regard (3) and (4) as one and the same state for the overall two-electron system. We would therefore expect to see the following distribution if measurements are made on the overall system: (1) with $\frac{1}{4}$ probability, (2) with $\frac{1}{4}$ probability, and the (3,4) conglomerate with $\frac{1}{2}$ probability. If such a distribution obtains in the long run, then the system is said to obey the Maxwell-Boltzmann statistics.

However, this is simply not what obtains in the laboratory. What we *do* observe is that (1), (2) and the (3,4) conglomerate state each occur with $\frac{1}{3}$ probability. The Maxwell-Boltzmann statistics leads us astray in this example! It turns out that electrons (more generally: fermions)

follow a different probability distribution altogether: the Fermi-Dirac statistics. The Maxwell-Boltzmann distribution provides a decent approximation of the system's behavior, but it is just that: an approximation. It is only in the classical regime that the Maxwell-Boltzmann distribution coincides with the Fermi-Dirac distribution.

What, then, accounts for the disparity between our intuitive probability distribution and the empirical evidence? Well, the problem seems to lie in the fact that we treated (3) and (4) as distinct states. When all states are equiprobable, and particular states are observed with – probability, there is indisputable evidence that we are dealing with exactly three, and not four, different states. We treated (3) and (4) as distinct because our intuitions tell us that although we cannot *detect* a difference, (3) and (4) still surely differ in virtue of the fact that A is in in (3) whereas it is in in (4). Similar things can be said for B. This line of reasoning turns on the crucial assumption that A and B are individuals. That is to say that they can be distinguished from their neighbors.

But electrons *cannot* be completely distinguished from each other, and these intuitions are merely artifacts of our everyday interactions with medium-sized objects. For example, if my friend and I each buy a copy of *Kind of Blue* from the record store, and if I break into his house that very night and swap my copy for his, although the switch is imperceptible to him, there is still a very real sense in which the overall distribution of copies of *Kind of Blue* has changed. But as the numbers show, this analogy is not a valid one to make because there is no real sense in which swapping A and B produces a different state. Our two electrons are *non-individuals*.

It is worth noting that we could run through this entire example again – and arrive at the same conclusion of non-individuality – if we investigated the *position* of particles A and B in a box partitioned into two halves. We would find that the state in which A is on the left half and B

is on the right half, and the state in which the two particles are exchanged, are really one and the same state. Again, the measured probabilities would lead us to this conclusion. The picture that emerges is again that A and B are non-individuals, and all that matters is the overall configuration of the particles in the box. The particles are no more than nameless faces in the crowd.

It is helpful to return to the bank account analogy in order to digest this example further. Imagine that, as is sometimes the case, two separate ATM's are located outside the bank. Suppose I have two paychecks to deposit, and I decide to deposit each paycheck into a different ATM. I can deposit paycheck A in the first ATM (and B in the second), or I can deposit paycheck B in the first (and A in the second). Whichever eventuality obtains, my bank account will be in the same overall state; the balance will have increased by the amount (paycheck A + paycheck B). The overall state of the bank account is indifferent to which ATM the checks were handled by, just as the overall state of the two-electron-system is indifferent as to which electron is in which state (assuming that *one* of them is in ψ_1 and *the other* is in ψ_2). This is all due to the fact that the dollars in the bank account are non-individuals: it makes no sense to speak of any particular dollar.

We are now ready to tie this example into LL. As non-individuals, these electrons cannot be considered in their own right. It only makes sense to speak of the entire two-electron system. The two electrons are indiscernible: although we are able to tell whether a given state is occupied (or doubly occupied), there is no way to tell one electron apart from the other. To reiterate Griffiths:

You can always tell the [classical] particles apart ... just paint one of them red and the other one blue, or stamp identification numbers on them, or hire private detectives to follow them around. But in quantum mechanics the situation is fundamentally different ... You can't paint an electron red, or pin a label on it,

and a detective's observations will inevitably and unpredictably alter its state, raising doubts as to whether the two had perhaps switched places.¹³

We have satisfied the antecedent of LL: the electrons are indiscernible from each other. But this only takes us halfway to forming a counterexample; we now need to deny the consequent. The argument that the two electrons are in fact *not* identical is both extremely subtle and extremely crucial to the overarching task of this essay. Thankfully, David Lewis articulates this point quite well:

Identity is utterly simple and unproblematic. Everything is identical to itself; nothing is identical to anything else except itself. There is never any problem about what makes something identical to itself; nothing can ever fail to be. And there is never any problem about what makes two things identical; two things can never be identical.¹⁴

In other words, by merely stipulating that there were *two* electrons, we guarantee their non-identity. Now it is often the case that two *names* pick out the same object – as in the case of ‘Clark Kent’ and ‘Superman’ – but this is not a case in which two objects stand in an identity relation to one another; there is just one object, and it can be called by one of two names. But to say that there are *two* objects (whether they be electrons, chairs, or dogs) precludes the very possibility of their being identical. Therefore, we have an example of two objects, A and B, which manage to be utterly indiscernible from each other yet fail to be identical. Conclusion: LL is false in this quantum context.

4.2 *Second Argument: The Gibbs Paradox*

We now shift our attention to a different example, this time in statistical mechanics. Consider a box filled with one particular type of gas molecule – molecules, say. We can

¹³ Griffiths (2005), p. 210.

¹⁴ Lewis (1986), pp. 192-193.

calculate the entropy of this system by counting the number of microstates – that is, by counting the number of ways that the molecules can be arranged throughout the box. Furthermore, we know from the 2nd law of thermodynamics that the entropy of this box will increase.¹⁵ If we consider the molecules to be individuals in calculating the entropy (that is, in counting the number of microstates), then we arrive at a very strange result known as The Gibbs Paradox. Calculating the entropy in this way, it would be possible to *decrease* the total entropy simply by inserting a partition in the middle of the box.¹⁶ This is problematic on two fronts. First, the addition or removal of a partition in a box of gas molecules should have *no effect whatsoever* on the total entropy.¹⁷ Second, even if we suspend our requirement that the entropy will not be changed, the fact that it decreases, rather than increases, is in strict opposition to the 2nd law.

So what is the resolution to this paradox? It turns out that in order to have the entropy remain unchanged during the addition or removal of a partition, we would need to divide the number of microstates that we counted by a factor of $N!$, where ‘ N ’ represents the number of molecules in the box. Furthermore, it turns out that this is precisely the factor which corrects for the “overcounting” of microstates.¹⁸ This tells us that if in counting the number of microstates we regard the exchange of two molecules as *not* yielding a different state – just as the two

¹⁵ More accurately: the entropy of a closed system will either stay the same or increase. *Most* accurately: it is *overwhelmingly improbable* that the entropy of a closed system will decrease.

¹⁶ On the mathematical level, this means that the entropy of a box of volume V will be *more* than the sum of the entropies of two boxes, each of volume $-V$.

¹⁷ See Schroeder (2000), p. 80 for an explanation of why this is so.

¹⁸ What do I mean here by ‘overcounting’? Suppose three friends want to watch a movie, but the only piece of furniture in the room is a loveseat, so only two of them can sit comfortably. Suppose further than sitting on the left or right half of the loveseat makes no difference; each person would be satisfied to merely sit on the sofa, because it means that they don’t have to watch the movie while sitting on the floor. How many ways are there of watching this movie? Given these constraints, it would seem that “three” is the proper response: either person A, person B, or person C must sit on the floor, and that makes three ways. One might also be tempted to answer with “six” because for each of the three aforementioned arrangements, the two on the loveseat could swap places, making six total arrangements. However, this would be an example of overcounting. As I *stipulated*, it makes no difference which side of the loveseat one sits on, so swapping the two people on the loveseat does *not* yield a new state. If one ignores this constraint and responds with “six,” he/she has overcounted.

couch-sitters swapping places does not yield a different state – then the entropy will remain unchanged during the addition or removal of a partition! The moral of the story is that we need to look at the Gibbs Paradox as a *reductio ad absurdum* on treating quanta as individuals. As Schroeder puts it, “[t]he best resolution of the paradox is simply to assume that all atoms [or molecules, or elementary particles] of a given type are truly indistinguishable.”¹⁹ But the fact that I stipulated that there were *many* molecules precludes the possibility of their being identical. These molecules are indiscernible – yet they are not identical (because there are *many* of them, not just one) – and constitute a counterexample to LL, just like the dollars in my bank account.

4.3 Third Argument: The Uncertainty Principle

The Heisenberg Uncertainty Principle places limitations on how well-defined observable quantities (position, momentum, spin, energy) can be. To take the most well-known example, consider the incompatible observables position and momentum. If the uncertainty in the position is high, then the uncertainty in momentum must be low, and vice versa.²⁰ As the joke goes, if Heisenberg were to be pulled over for speeding – having been clocked in at exactly 70mph, say – then he would have to admit that at the time his speed was being measured he was completely lost.

The orthodox interpretation of quantum mechanics takes the Uncertainty Principle to be a metaphysical rather than epistemological constraint. It says nothing about our knowledge of the values of observables, and it says nothing about the accuracy of the equipment in our

¹⁹ *Introduction to Thermal Physics* (Addison Wesley Longman, 2000), p. 81.

²⁰ I am not using ‘uncertainty’ in an epistemological sense here, but rather in a statistical sense. The uncertainty in position can be thought of as the standard deviation that is calculated from making a large number of position measurements on an ensemble of identically prepared systems. Likewise for the uncertainty in momentum. The statistical notion of standard deviation is a brilliant way to quantify an epistemological notion like confidence (or lack thereof); this explains the physicist’s choice in terminology. But when we use the word ‘uncertainty’ in experimental contexts, it is meant to invoke the statistical connotation only.

laboratories. The principle is a direct result of interpreting the wavefunction as a measure of the probability of finding the particle in a particular location. The derivation of this principle makes no reference to humans or laboratory equipment; it refers only to the wavefunction itself.

Let us now imagine that our task is to perform an act of discernment between electrons. We can either appeal to the spatiotemporal properties of each electron (those that vary from situation to situation, such as position or momentum) or the non-spatiotemporal properties (those that remain constant from situation to situation, such as mass or charge). I take this list to be exhaustive: all electron properties can either be classified as spatiotemporal or non-spatiotemporal.

The non-spatiotemporal properties will not help us in our task of discernment, for all electrons share exactly the same mass, charge, etc. Using the non-spatiotemporal properties to discern electrons is as nonsensical as trying to sort a group of siblings on the basis of who their parents are; the siblings all have the same parents, and the electrons all have the same mass! The spatiotemporal properties are of no help either, for the proper interpretation of the Uncertainty Principle tells us that quantum particles simply do not have well defined positions, momenta, energies, etc. until an act of measurement intervenes and collapses the wavefunction. Using the spatiotemporal properties to discern two electrons is as nonsensical as trying to sort a group of black-and-white photographs on the basis of color; those properties just aren't there! If these two sets of properties do in fact form an exhaustive list, then there simply does not exist any way to discern one electron from another. Plain and simple. But the fact that I stipulated that there were *two* electrons precludes the possibility of their being identical. Thus LL fails in this context as well.

4.4 Fourth Argument: The Pauli Exclusion Principle

Imagine that we are to construct a two-particle system out of two electrons, A and B, and two states, ψ_1 and ψ_2 , that the electrons can occupy. It is therefore possible that A is in ψ_1 and B is in ψ_2 , in which case the overall wavefunction will simply be the product $\psi_1(A)\psi_2(B)$. But it is also possible that A is in ψ_2 and B is in ψ_1 , in which case the overall wave function will be the product $\psi_2(A)\psi_1(B)$. Notice how A and B have exchanged states compared to a moment ago. Now if we are to take the indiscernibility of quantum particles seriously, then when constructing the wavefunction for the overall system we must remain noncommittal as to which particle is in which state. We do this by taking a *linear combination* of each of these two possibilities. The two-electron system is therefore in the state $C[\psi_1(A)\psi_2(B) \pm \psi_2(A)\psi_1(B)]$ where C is just a normalization constant. When we invoke the plus sign above, then we say that the particles in question are bosons; when we invoke the minus sign, we say they are fermions.

What if ψ_1 and ψ_2 were to be the *same* state? In other words, what if $\psi_1 = \psi_2$? This would mean that the two terms, $\psi_1(A)\psi_1(B)$ and $\psi_1(B)\psi_1(A)$, would in fact be equal to one another. Then for bosons, the particles that choose the plus sign, we have:

This is unproblematic. But for fermions, who choose the minus sign, we have:

This would mean that the overall wavefunction is zero, which is tantamount to saying that the two-particle system doesn't even exist! So if the two-fermion system is to in fact exist, then the two fermions must be in different quantum states. This is the famous Pauli Exclusion Principle (PEP). “It is not (as you may have been led to believe) a weird ad hoc assumption, but rather a

[mathematical] consequence of the rules for constructing two-particle wavefunctions” for fermions.²¹

Weyl has suggested that the PEP can be construed as a *vindication* of LL, and that we should go so far as to use the phrase ‘Pauli-Leibniz Principle.’²² However, I believe that we ought to draw the opposite conclusion: PEP *refutes* the validity of LL in the case of fermions. If we massage PEP into the form of a conditional, then it states the following.

PEP: If two fermions (of the same type) occupy exactly the same quantum state, then the two-particle system would not even exist in the first place.

Compare with LL applied to two fermions (of the same type).

LL: If two fermions (of the same type) are indiscernible, then they are in fact identical.

I take it that *occupying exactly the same quantum state* and *being indiscernible from each other* are just two ways of stating the same thing – the former in the language of physics, the latter in the language of metaphysics. But from the same antecedent, these two conditionals allow us to draw drastically different conclusions. Whereas PEP tells us that true indiscernibility implies the nonexistence (impossibility) of the two-fermion system, LL tells us that true indiscernibility implies an identity relation among the fermions. These two conclusions are clearly inconsistent with each other: in order for an identity relation to hold (LL assures us that one does) the system about which we are speaking must exist in the first place (but PEP assures us that it does not).

PEP is derived as a mathematically necessary result. Furthermore, experimental confirmation of its universality increases our confidence in its truth. And since PEP draws conclusions that are inconsistent with LL, we must conclude that LL fails, if only in the context of two-fermion systems.

²¹ Griffiths (2005), p.204.

²² Weyl (1950).

4.5 *Haecceities to the Rescue?*

These four examples have shown us that major anomalies result from treating elementary particles (or atoms of a given kind, or molecules of a given kind) as individual entities. I now want further clarify this point – that such quanta are non-individuals – couched in terms of haecceity. Teller articulates this notion quite well:

If your necktie and mine are qualitatively identical, many people nonetheless feel that there is a “this one” and a “that one” about them ... independent of all their distinguishing properties and relations. The world would have been different ... if I had put on your tie and you had put on mine, indeed if your tie had gone through my tie’s entire history and my tie had gone through yours.²³

The idea is that each object possesses a primitive “thisness” – the property of being that object – something independent of all its traditional properties and relations, which sets it apart from all other objects. A thing’s haecceity is “some inscrutable aspect, unlike anything we ordinarily think of as a property, which a thing must be supposed to have to make sense of its identity and distinctness from other things.”²⁴ In the case of macroscopic objects, the notion of haecceity may seem a bit unnecessary. If the job is to individuate one object from a group of qualitatively indiscernible counterparts (e.g., a blade of grass in a field), then why not appeal to spatiotemporal properties?

For one thing, although spatiotemporal properties *would* serve to individuate a blade of grass in a field, the blade presumably has its position accidentally, whereas its haecceity is a matter of essence. I take it that no object has its position in spacetime *essentially*, and therefore that spatiotemporal properties do not have the permanence, from situation to situation, that haecceities do when it comes to the task of individuation. Secondly, and more importantly to the

²³ Teller (1998), p. 117.

²⁴ *Ibid.*, p. 117.

subject matter at hand, spatiotemporal properties generally fail to be well defined on the quantum scale via the Uncertainty Principle; strictly speaking, electrons don't *have* position!

As we have already seen it is hopeless to try to appeal to an electron's position, momentum, etc. in order to separate it from the crowd. And it is of no hope appealing to any of the electron's non-spatiotemporal properties – e.g., mass, charge – for these are, of course, exactly the same for all electrons! So if spatiotemporal properties cannot do the work, and neither can non-spatiotemporal properties, then how are we ever going to individuate among electrons? It seems that we must posit a haecceity – a pure “thisness” which transcends any normal conception of property or relation – to do the work of individuation.

Although I will not delve into arguments against the notion of haecceity here, it is clear that such arguments would provide one additional route to the conclusion that quanta are non-individuals. Haecceities seem to be the only way to preserve quantum individuality in light of the fact that both spatiotemporal and non-spatiotemporal properties fail pack any individuating power on the quantum level. An anti-realist regarding haecceities has not a single tool of quantum individuation at his disposal, and is therefore left with the conclusion that quantum particles are non-individuals.

4.6 Max Black's Iron Spheres: A Neglected Insight

In Max Black's famous paper “The Identity of Indiscernibles,”²⁵ interlocutor B launches a powerful attack against A, a defender of the principle of the identity of indiscernibles. Although Black's thought experiments represent valiant efforts in the business of producing counterexamples to LL, it is my belief that they ultimately fail. In this section, we will articulate

²⁵ Black (1952).

the reason why they all fail, and subsequently, why the two-electron system mentioned in 4.1 (First Argument: Quantum Statistics) succeeds in producing a counterexample to LL.

Let us imagine Black's universe which contains nothing but two qualitatively identical iron spheres, one mile in diameter each, which rest in each other's vicinity. The idea is that "every quality and relational characteristic of the one would also be a property of the other."²⁶ For example, each will have the relational property of being at a certain distance from a sphere one mile in diameter. A and B bicker for a number of pages regarding our ability to name and refer to these spheres. If we could refer to these spheres properly, then we could say that Castor (the name Black chooses) has the property of being two miles away from Pollux, whereas Pollux has the property of being two miles away from Castor. But in order to name them (and, subsequently, in order to nail them down on coordinate axes) B says, we must first introduce a third object – the space traveler/observer – which violates B's original stipulations. In other words, spatio-temporal relations are not available to serve as individuator until someone to measure these spatio-temporal relations arrives.

This point is quite contentious, and it is with this point that Black's thought experiment ultimately fails to be a bulletproof counterexample to LL. It is not immediately clear whether or not spatial properties like those just described are dependent upon the presence of an observer. Personally, my sympathies lie with A: "You surely don't want to say that the arrival of the name-giving traveler creates spatial properties? Perhaps we can't name your spheres and therefore can't name the corresponding properties; but the properties must be there."²⁷ However, it is unimportant what the answer to this question is. By slightly tweaking Black's thought

²⁶ Black (1952), p. 99.

²⁷ Black (1952), p. 101.

experiment, we can completely sidestep the contentious status of spatial-properties-sans-observer, while still retaining its overall spirit!

The key lies in the shift from macroscopic to microscopic objects. Instead of imagining two gigantic metal spheres, let us instead imagine that we are dealing with two electrons. In this case, it is futile to debate over whether the objects have their spatial properties in the absence of an observer: observer or not, the two electrons simply do not have well defined positions! In fact, in the case where the momenta of the two electrons are precisely known, their positions are so poorly defined that it makes most sense to say that they don't even have any such property as position.

Let us take stock of everything now. The two electrons trivially share all their non-spatiotemporal properties (charge, mass, etc.) This is what Black aimed to accomplish by stipulating that the spheres were "qualitatively identical." Additionally, according to the orthodox interpretation of quantum mechanics, the electrons do not even *have* any non-spatiotemporal properties. Now we can just run the argument from 4.3 (Third Argument: The Uncertainty Principle): if the non-spatiotemporal properties are not up to the task of individuation, and the spatiotemporal properties don't even exist, then there is no conceivable way in which the two electrons can be distinguished. They are therefore indiscernible. However, Black has stipulated that there are *two* of them, so they cannot be identical. Thus, we have countered LL.

Black actually mentions this very example, but fails to see the benefits of the shift from iron spheres to electrons: "Don't physicists say something like this about the electrons inside an atom? We can verify *that* there are two, that is to say a certain property of the whole configuration, even though there is no way of detecting any character that uniquely characterizes

any element of the configuration.”²⁸ This neglected insight, if properly attended to, would have allowed him to produce a much stronger counterexample to LL.

²⁸ Black (1952), p. 103.

Metaphysical Implications of Non-individuality

5.1 Entanglement

As we learned in the previous section, if we take the formalism of quantum mechanics seriously, then there is no meaningful sense in which quanta can be regarded as individuals. In this section we will connect this notion of non-individuality with the phenomenon of quantum entanglement.

Two particles enter into an entangled state when “the two-particle state cannot be expressed as the product of two one-particle states, and for which, therefore, one cannot really speak of “the state” of either particle separately.”²⁹ It is clear that such a situation would suggest a form of holism that is “surprisingly ... absent from classical physics.”³⁰ For example, the famous Einstein-Podolsky-Rosen paradox deals with a spin-zero particle which decays into two spin- $\frac{1}{2}$. These two quanta are *entangled* in a manner wholly unfamiliar to the realm of macroscopic objects. But the phenomenon of entanglement is by no means restricted to specific decays patterns. Entanglement also occurs when two electrons occupy what is known at the singlet state. Furthermore, the two electrons in a Hydrogen molecule in its ground state *necessarily* occupy the singlet state; thus, entanglement is ubiquitous.

It seems to me that the non-individuality thesis (that quanta cannot be investigated separately from other quanta of the same type) and the entanglement thesis (that certain two-particle state cannot be broken down into combinations of one-particle states) are both gesturing toward the holistic approach that Maudlin describes. These are simply two different routes to the

²⁹ Griffiths (2005), p. 422.

³⁰ Maudlin (1998), p. 46.

same conclusion. In this section we will investigate the metaphysical implications of the holistic attitude that seems to be inseparable from quantum theory.

5.2 *Reductionism & Holism*

According to Maudlin, the aim of modern physics is, by and large, to understand wholes in terms of their parts.³¹ The temperature of a glass of water is to be understood as the mean kinetic energy of the molecules; the magnetization of an iron bar magnet is to be understood as the summation of the magnetic moments of each individual atom; the complex interworking of the human body is to be understood in terms of bodily systems, which themselves are to be understood in terms of organs, which are to be understood in terms of cells, etc. The macroscopic properties (the room's temperature, the bar magnet's magnetization, the body's overall state) are said to *supervene* on the microscopic properties (atomic motion, magnetic moments, the states of cells). No change in the macro-properties without a corresponding change in the micro-properties.

However, if we dig deep enough into the supervenience base we will eventually be presented with a whole (the singlet state, for example) which simply cannot be decomposed into parts (two separate one-particle states). This seems to be an insuperable obstacle for the reductionist enterprise, and ironically enough, one that is necessitated by the orthodox interpretation of our most powerful and experimentally verified physical theory. What I want to suggest in this section is that the holism required by quantum theory is not *detrimental* to our precious reductionist mindset.

Consider Newton's theory of gravitation. According to this theory, each particle in the universe exerts a force on every other particle that is proportional to the product of the masses of

³¹ Maudlin (1998), p. 46.

the two particles and inversely proportional to the square of the distance between them. Each particle in the universe is therefore “connected” to the rest insofar as it produces a gravitational tug on every other particle. The tug that my body feels from some particular Hydrogen atom in Alpha Centauri is certainly of no practical or measurable importance, but nevertheless still exists in principle. Thus, a sense of holism emerges from Newtonian gravitation in so far as a *complete* description of the gravitational influences on a given particle must take into account matter from across the galaxy. But in literally all situations that we encounter, these distant effects are of no importance; the distance between my body and that Hydrogen atom in Alpha Centauri is so large as to make the gravitational force between us essentially zero. Newtonian holism is quite tolerable, and one who recognizes the gravitational “connection” between each particle in the universe should in no way feel that his reductionist program has been threatened: this type of holism is for the most part negligible.

Much in the same way, I think that quantum holism is not detrimental to the reductionist enterprise. The holistic attitude does dictate that “in principle every electron in the universe is linked to every other one,”³² meaning that a *complete* description of the state of an electron in my laboratory must mention some particular electron in Alpha Centauri, say. But such a description is of course astronomically pedantic, and physicists are keen to recognize that such a connection is negligible for all practical purposes. If this connection truly mattered, then “you wouldn’t be able to talk about any [electron] unless you were prepared to deal with them *all*.”³³ We clearly do succeed in talking about lone electrons³⁴ so we can only conclude that the connection is in most cases negligible.

³² Griffiths (2005), p. 209.

³³ Griffiths (2005), p. 209.

³⁴ We succeed in the case just described. We would obviously fail, however, in the case of the singlet state.

The same idea – that holistic factors exist in *principle*, but are in most cases negligible – also surfaces on other, more commonplace, levels. As another example, consider the case of my identical-twin aunts, Cecelia and Roselee.³⁵ Both have been intricate parts of each other’s lives for over 60 years. They have lived and worked together for decades, they share similar opinions on almost any issue, and as a result have formed a deep bond which bears heavily on their day-to-day lives. This happens to be an actual example, but we can of course imagine other twins whose involvement in each other’s lives is infinitely more intricate than in the case of my aunts. But to focus on the example at hand, it is fair to say that a truly complete description of Cecelia’s psychological state *must* mention Roselee, and vice-versa. However, a truly complete description is not always what we are after, and an incomplete description will often suffice. For example, if we wish to understand why Cecelia is in a psychological state of glee and excitement, the fact that her lotto numbers were just picked suffices to provide such an explanation. Perhaps Cecelia is partially excited because she can now afford to take a vacation with her sister, but this information only helps to clarify her excitement, and is in no way necessary to a satisfactory explanation of her current psychological state. Analogously, the holism brought forth by quantum mechanics can often be neglected; we can provide satisfactory explanations of quantum phenomena by simply ignoring it.

The way in which one reacts to the holism necessitated by quantum mechanics falls along a spectrum. On one end we find the knee-jerk reaction that this holism is in some way detrimental to the task of science. But this is not the proper response. As I have argued, our normal reasoning practices already admit a form a holism. If we accept that a *satisfactory* description of Cecelia’s psychological state needn’t mention Roselee (although a *complete*

³⁵ Actually, the same point can be made with *fraternal* twins. The parallel with indistinguishable quanta will be more transparent for *identical* twins, however.

description will mention her), then we should have no qualms about giving a merely satisfactory description of the state of a lone quantum particle.

But the other end of the spectrum is equally troubling; we should not *embrace* the holism as new-age spiritualists certainly feel compelled to do. Maudlin articulates this quite well:

...we should note that this form of holism has no general moral or methodological consequences. The defenders of the so-called holistic health movement, for example, can draw no comfort from quantum theory. The “holism” advocated there has to do with the idea that humans are very complex organisms, and that the biological state of any part may be influenced by a myriad of factors, including diet, exercise, and mental state ... Nor does the holism of quantum mechanics imply an ethically charged view of the universe, according to which all things are one, or are morally interconnected. If quantum mechanics inspires people to become vegetarians, or strengthens their resolve to promote world peace, then the effect is merely via a vague metaphor.³⁶

The task at hand, then, is to find a comfortably middle ground between a dogmatic distaste for quantum holism, and an overly sentimental and spiritualistic embrace of it. However, that is a considerable task – one for another essay.

³⁶ Maudlin (1998), p. 55.

6

Conclusion

The task undertaken in this essay was to tell a consistent story about the behavior of quantum particles. From their bizarre behavior emerges the conclusion that such particles do not have individuality in any meaningful sense of the word; they are non-individuals. However, the paradigmatic cases of non-individuality (illustrated by the four arguments in section 4) are not ones in which the *two* particles in question were *identical* to each other. As the helpful quote from Lewis urges, we should not need more than a moment's thought to reach that conclusion: "there is never any problem about what makes two things identical; two things can never be identical."³⁷ Thus, these examples show that LL fails in certain quantum contexts. But that is not to say that LL is false in *any* context; a counterexample to LL with regard to macroscopic objects has yet to be produced. However, many philosophers attribute a *universal* range of applicability to this principle. And in light of the examples I have provided, it is clear that Leibniz's Law needs to be construed much more modestly than this if it is to be construed at all.

³⁷ Lewis (1986), pp. 192-193.

7

References

- Castellani, Elena (1998), *Interpreting Bodies*, Princeton University Press.
- Feynman, Richard (1965), *The Character of Physical Law*, BBC.
- French, Steven and Krause, Decio (2006), *Identity in Physics*, Clarendon Press-Oxford.
- Griffiths, David (2005), *Introduction to Quantum Mechanics*, Pearson Education.
- Maudlin, Tim (1998), "Part and Whole in Quantum Mechanics," in Castellani 1998, pp. 46-60.
- Quine, Willard van Orman (1949), "Identity," Theodore and Grace de Laguna Lecture (Transcript).
- Russell, Bertrand (1912), *The Problems of Philosophy*, New York: Barnes & Noble.
- Schroeder, Daniel (2000), *An Introduction to Thermal Physics*, Addison Wesley Longman.
- Teller, Paul (1983), "Quantum Mechanics, the Identity of Indiscernibles, and Some Unanswered Questions," *Philosophy of Science* 50 (2): 309-319.
- (1998), "Quantum Mechanics and Haecceities," in Castellani 1998, pp.114-141.
- Weyl, Hermann (1950), *The Theory of Groups and Quantum Mechanics*, New York: Dover.